Configuring Oracle10g Automatic Storage Management

Single Node Installation on Windows 2000

Authors: Michael Hallas, Robert Smyth and Philip Newlan

Contact details: Mike.Hallas@oracle.com

Robert.Smyth@oracle.com

Philip.Newlan@oracle.com

Team: Oracle Technology Solutions Group, Oracle (UK)

Date: 24th December 2003

Version: Draft 1E

Software: Oracle Database 10*g* Release 1 (Beta)

IMPORTANT

THIS DOCUMENT IS BASED ON BETA CODE.

FUNCTIONALITY MAY CHANGE IN THE PRODUCTION RELEASE.

THE AUTHORS CAN NOT OFFER ANY ASSISTANCE WITH THE CONTENTS OF THIS DOCUMENT

Contents

1. Introduction

This paper walks through the creation of Oracle's Automatic Storage Management (ASM) using a laptop running Windows 2000.

NOTE: This is not a proposal for a production architecture! The aim is to create a "sandbox" to demonstrate the main features of ASM and to learn the new techniques.

ASM uses the Stripe And Mirror Everything technique developed to ease the management of Oracle databases. Mirroring ensures that no data is lost and Striping ensures that all storage components are used equally with no hot spots.

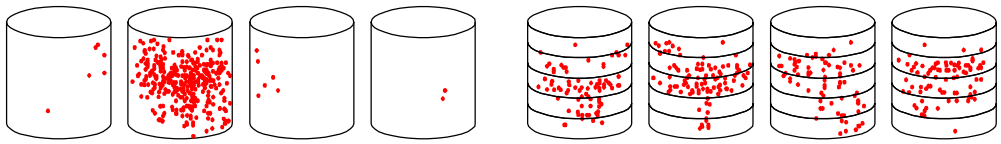

Unstriped Disks **Striped Disks**

Once implemented ASM uses disk groups and allows tablespaces and so tables to be created upon these disk groups. This removes the need for Logical Volumes, File Systems and Files to be managed thus dramatically reducing management requirements.

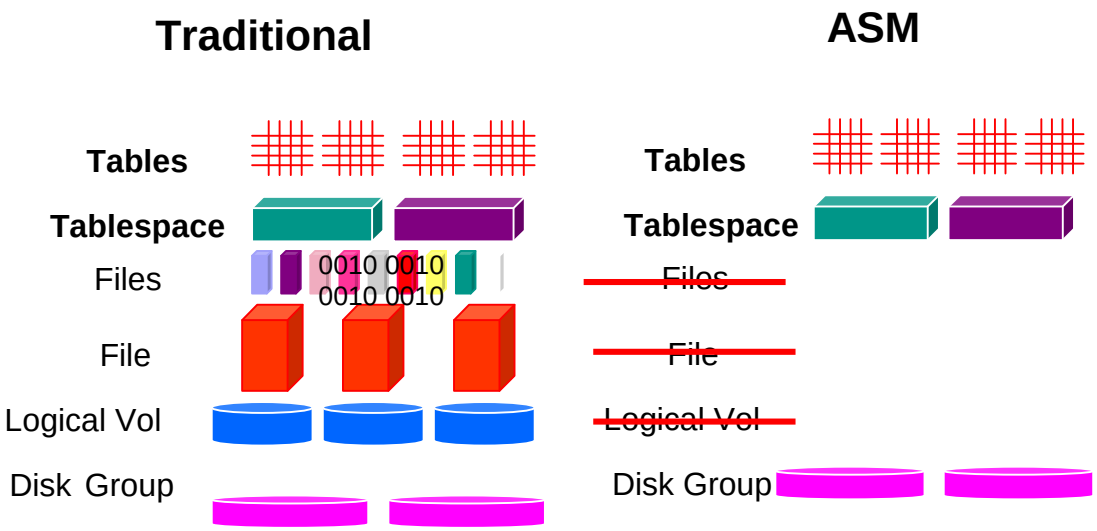

This paper outlines the pre-requisites needed for ASM, installs ASM and covers two hands on sessions to show the benefits of ASM.

2. Creating Logical Drives

2.1 Background to creation of Logical Drives

In order to implement Automatic Storage Management (ASM), we need to create logical drives that are unformatted and which ASM can use to create disk groups upon which the database tablespaces may be placed.

NOTE: In Windows, the standard file API cannot read a partition directly, so for ASM to access the disk we use a label – in this case this is the drive letter.

A step-by-step method for this is included below using Windows Disk Manager against Free Space on the laptop's hard drive.

There are other methods suitable to do this. If you prefer feel free to try a different method. Other methods that spring to mind include running `crlogdr` from the `preinstall_rac` folder of Disk 1 of the Oracle9*i* Real Application Clusters software distribution or using Partition Magic.

In our example, we have twenty Gbytes of free space on the hard disk and we will create eight 1Gbyte logical drives, unformatted but with drive letters.

2.2 Step by Step Method of creating Logical Drives using Windows Disk Manager

To run the Windows Disk Manager applet, right click on `My Computer`, click `Manage` and click `Disk Management`

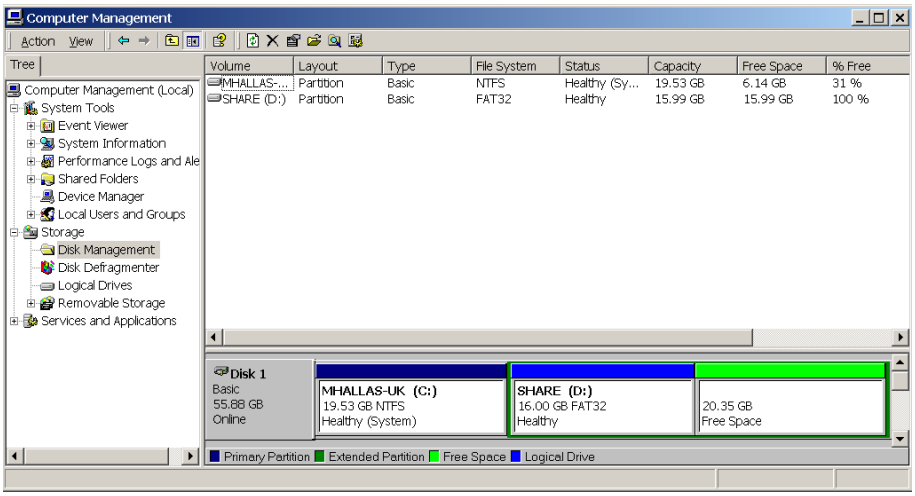

We will now create additional logical drives using storage allocated from the free space of your laptop's hard drive. Right Click on the `Free Space` and click `Create Partition`

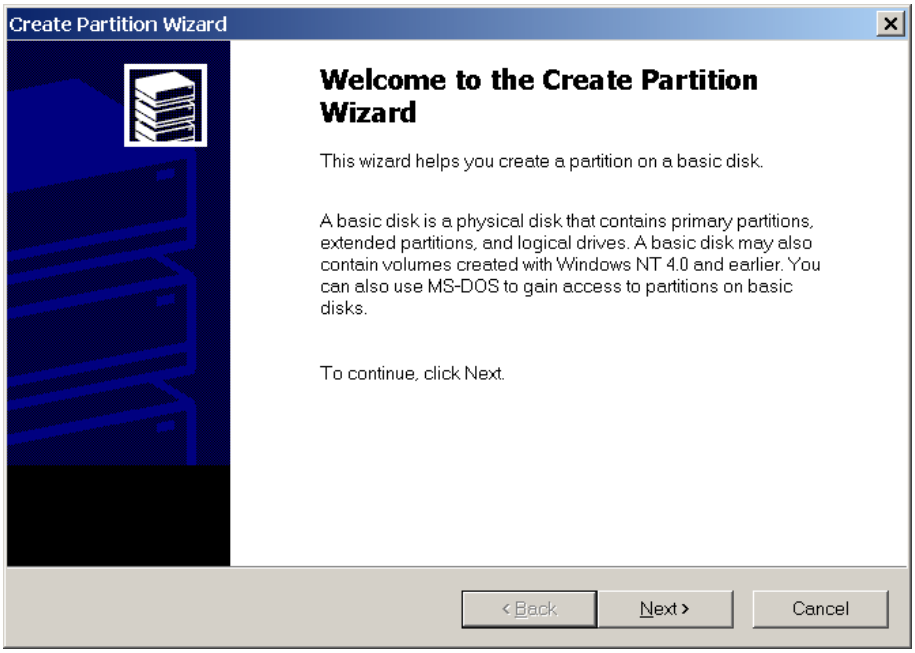

Click "Next"

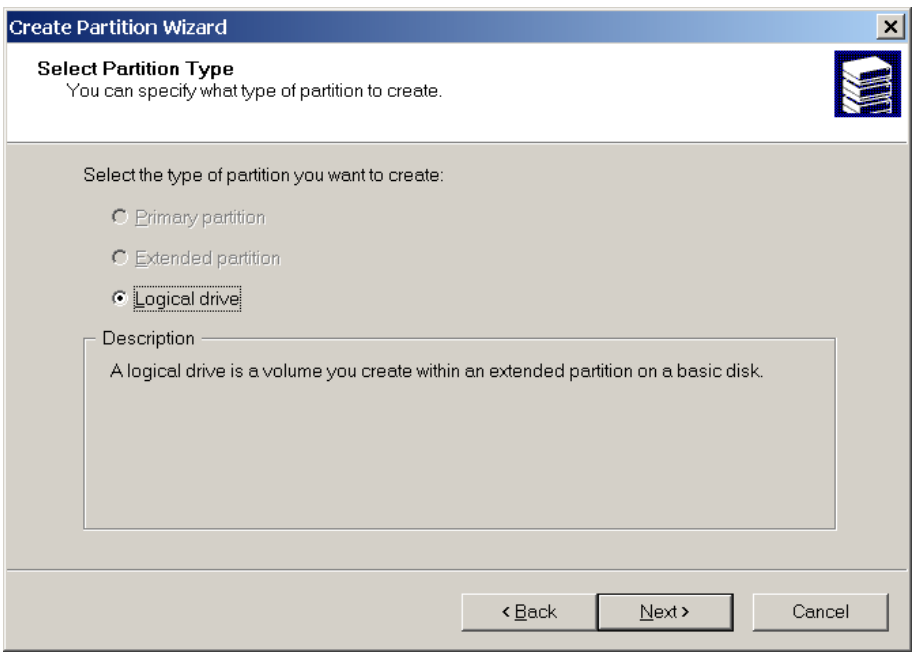

Click "Next"

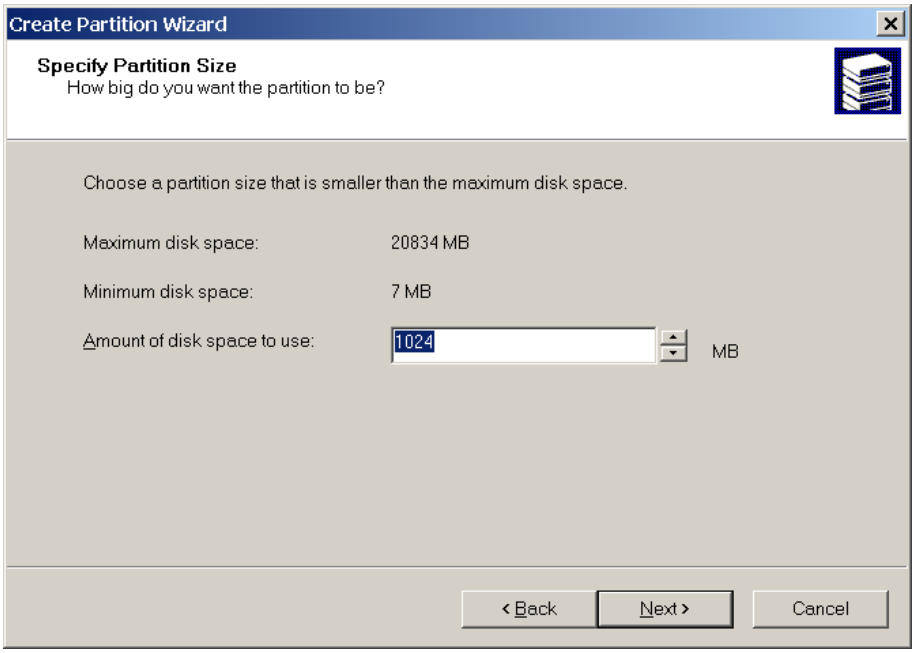

Enter 1024 in the "Amount of disk space to use" and Click "Next"

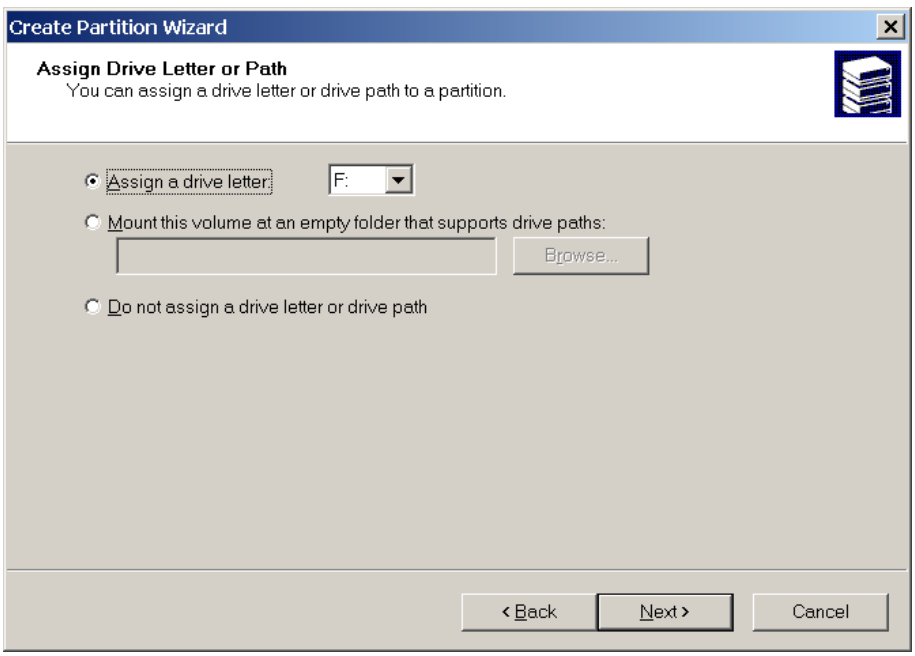

Click "Next"

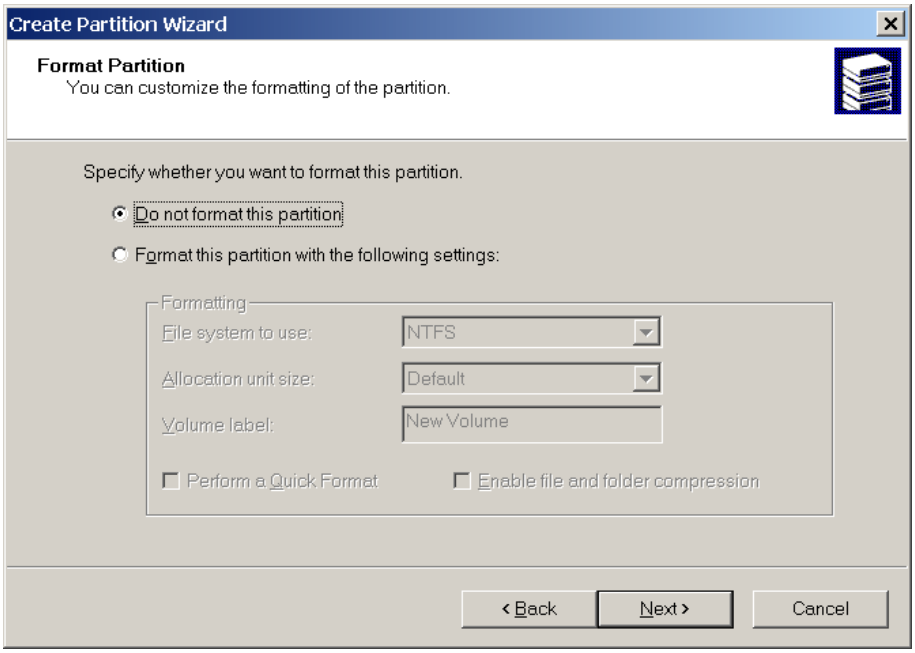

Select "`Do not format this partition`" and Click "`Next`"

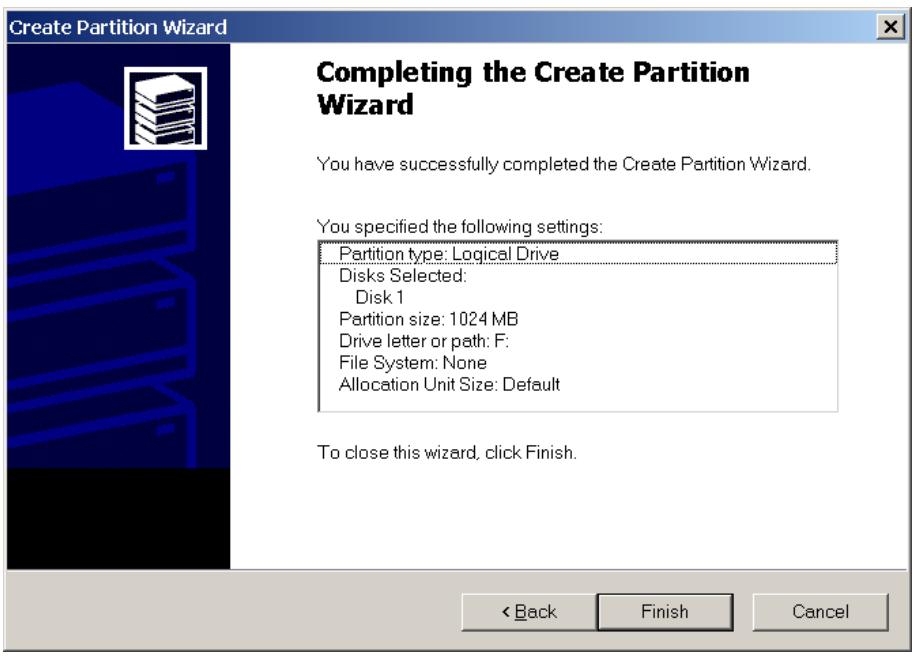

Click `Finish`

Repeat "`Create Partition`" until you have eight 1Gbyte logical drives. Note that you may have to reboot if the Windows Disk Manager requests this. Your final layout should look along the lines of the following illustration.

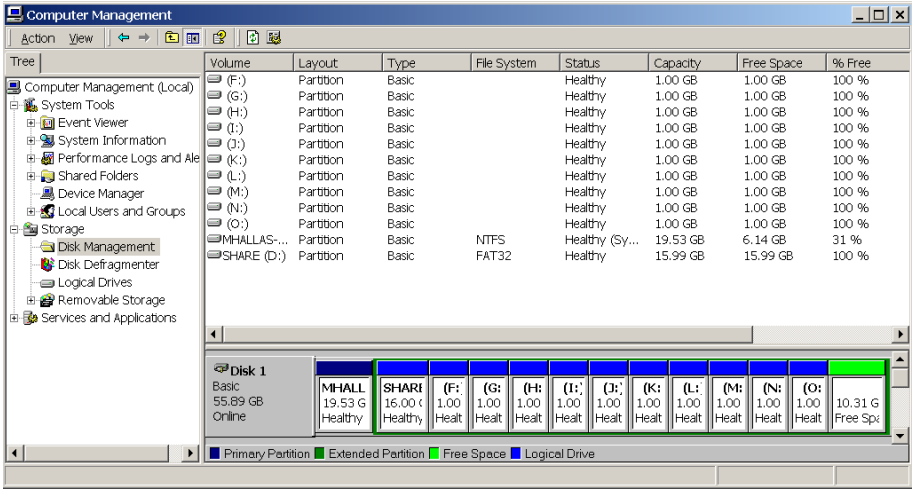

Once complete, you should have eight 1Gbyte logical drives labelled `F:` through to `M:`, with no file system. The above illustration shows two additional 1 Gbyte drives labelled `N:` and `O:` created during the later hands on sessions.

3. Creating a database running ASM

3.1 Pre-requisites

In the next section we will use the Oracle Database Configuration Assistant (DBCA) to create a database running Automatic Storage Management (ASM). DBCA gives three choices for storage: traditional file system, raw or ASM. On selecting ASM we will create a database disk group containing two failure groups with 3 logical drives in each (this shows Oracle ASM mirroring) and 2 external groups for the redo logs (showing that we can still use external mirroring technology).

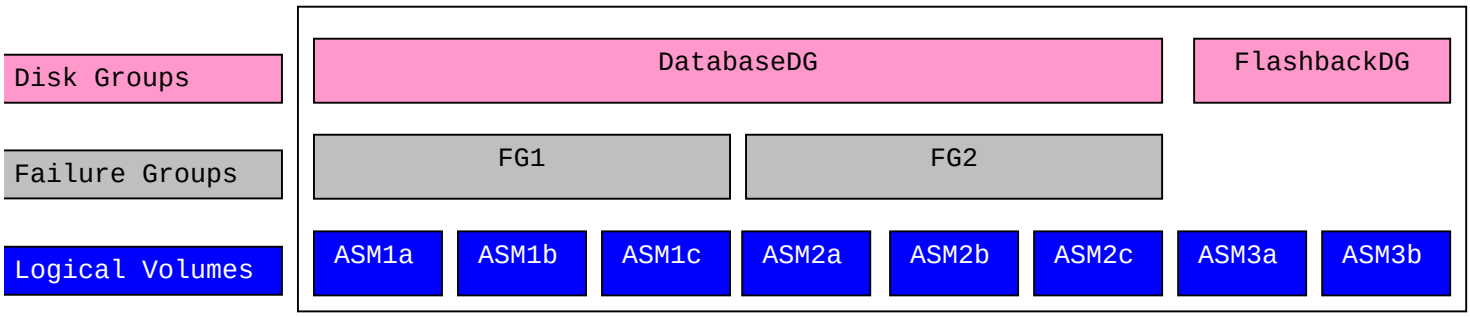

In the Windows implementation below, we will map the Logical Volumes to the eight 1Gbyte drives `F:` to `M:` created earlier.

NOTE: Prior to running DBCA you need to ensure that `OracleCSService` service is running.

3.2 Creating the Database

Launch DBCA from the Start Menu.

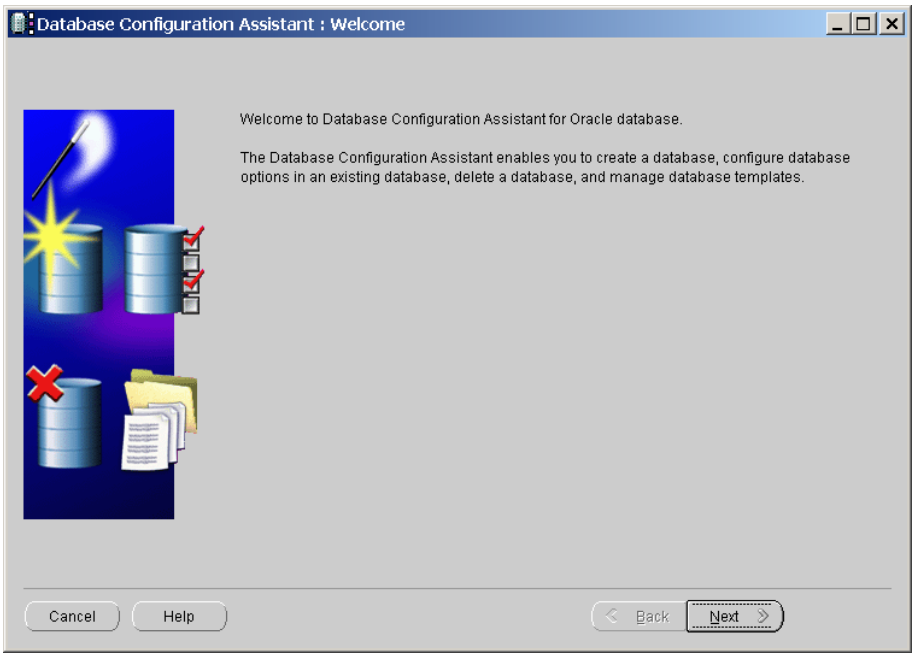

Click "`Next`"

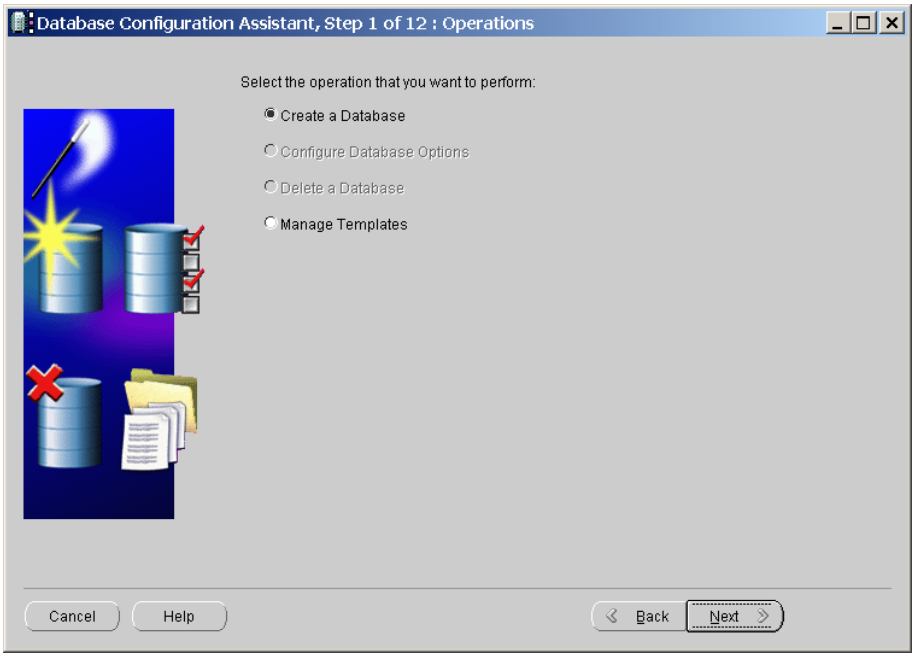

Click "Next"

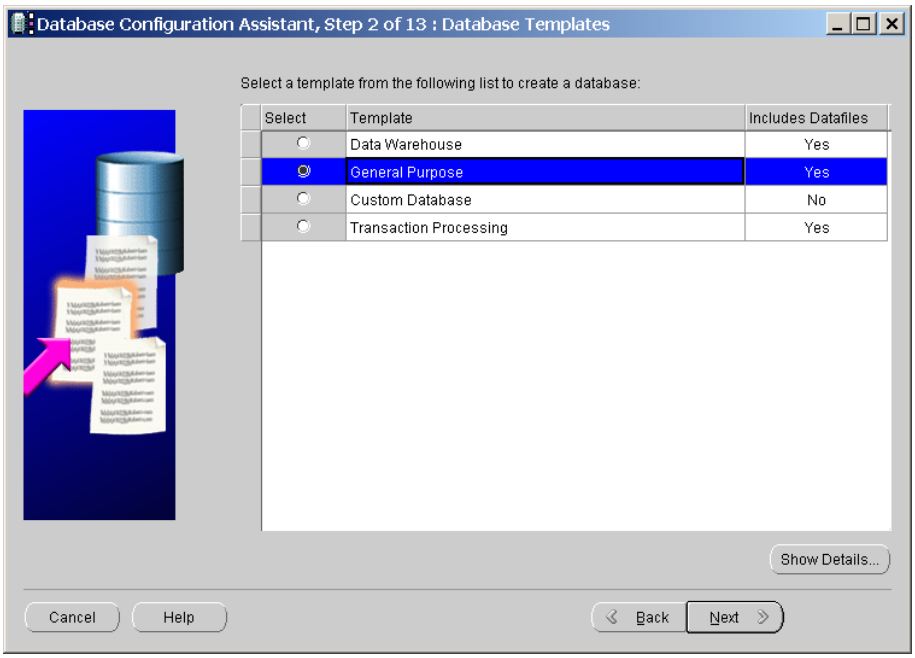

Select the "General Purpose" template and click "Next"

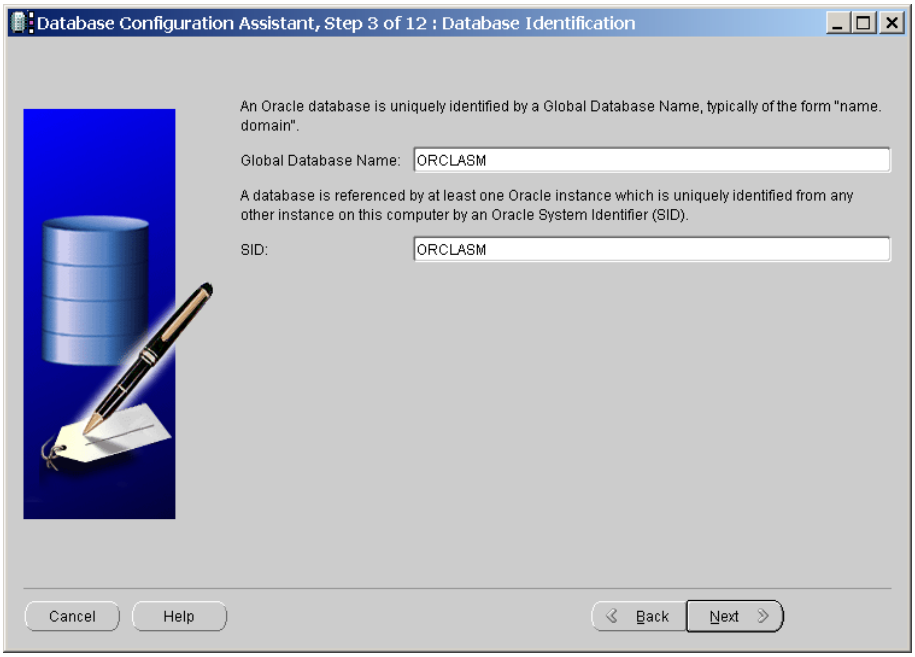

Enter your `Global Database Name` and click "`Next`". In our example we have called the database Oracle – `ORCL` Automatic Storage Management – `ASM`.

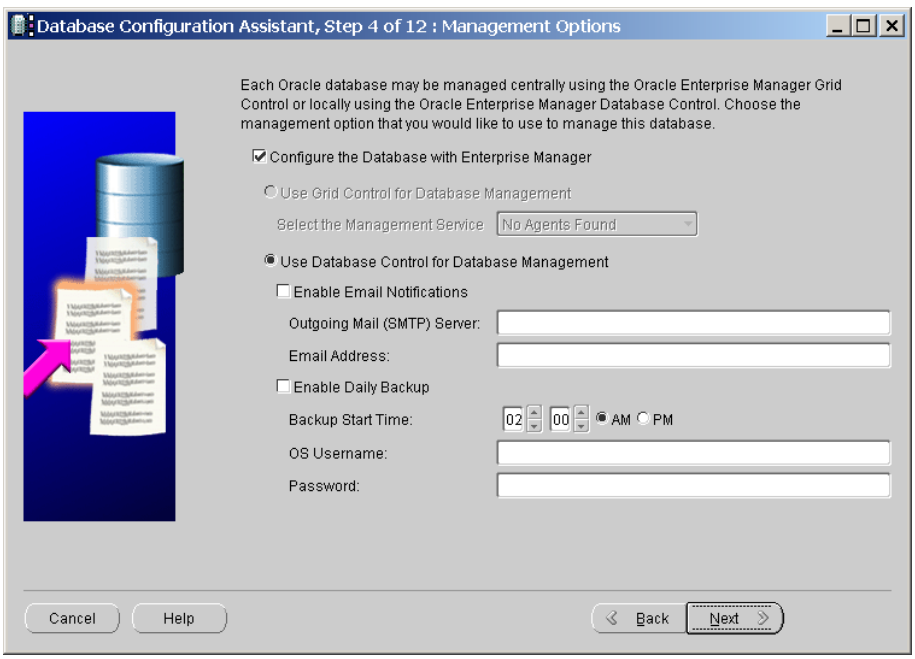

Ensure that "`Configure the Database with Enterprise Manager`" is checked and click "`Next`"

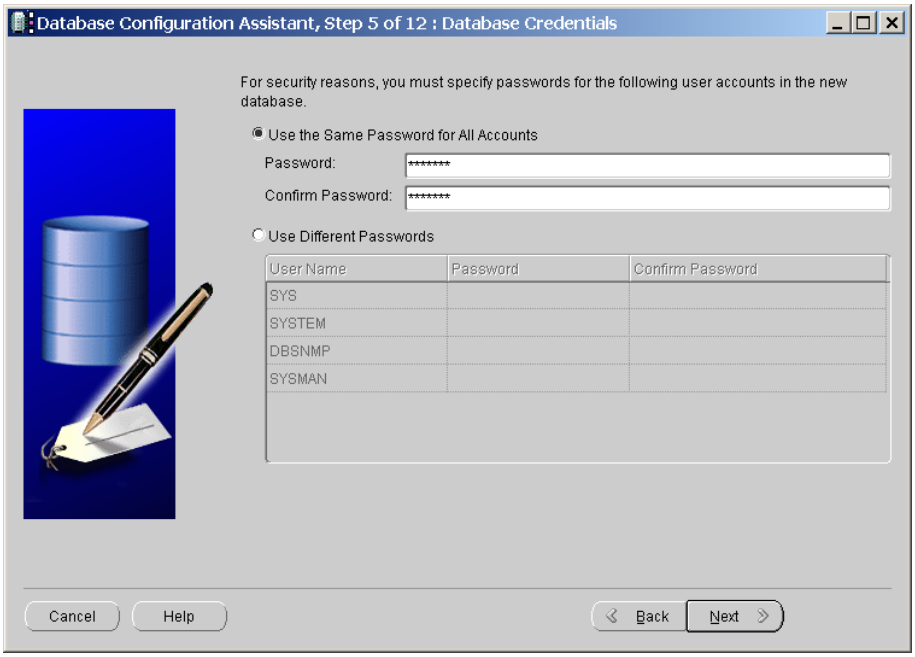

Select "`Use the Same Password for All Accounts`" and enter a password of your choice and click "`Next`"

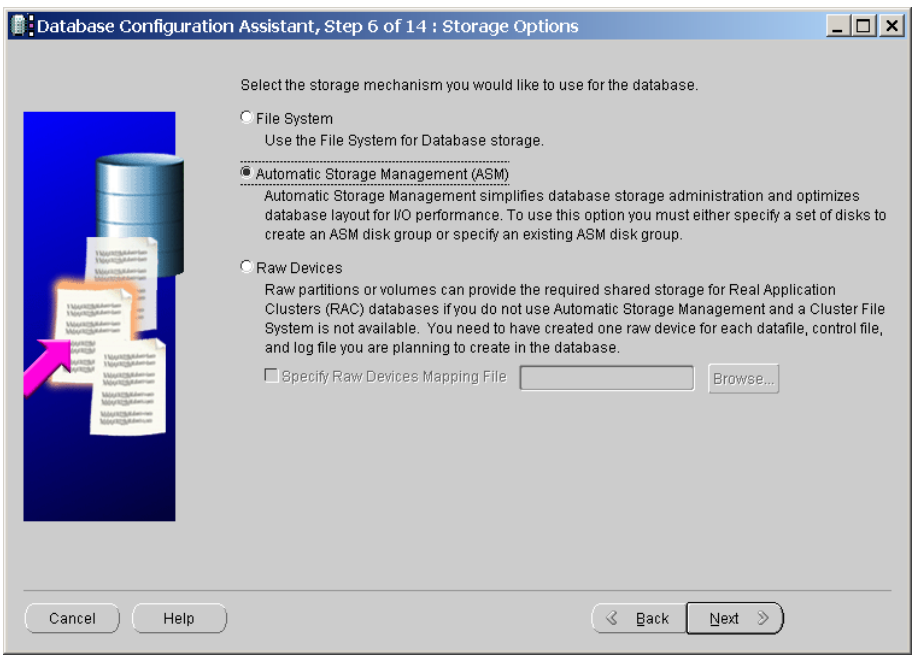

We wish to use Automatic Storage Management so select "`Automatic Storage Management (ASM)`" and click "`Next`".

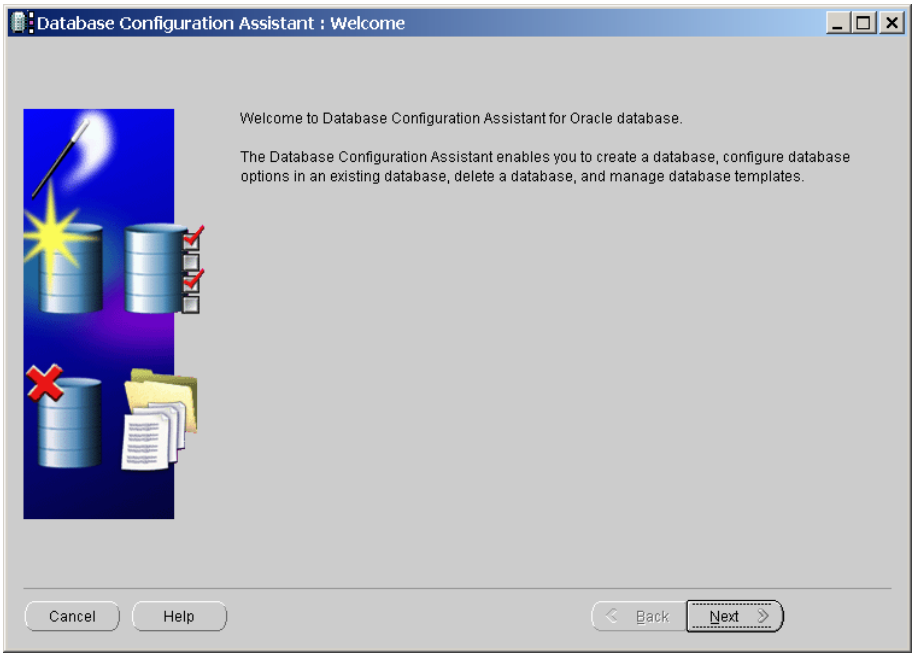

Click "Next"

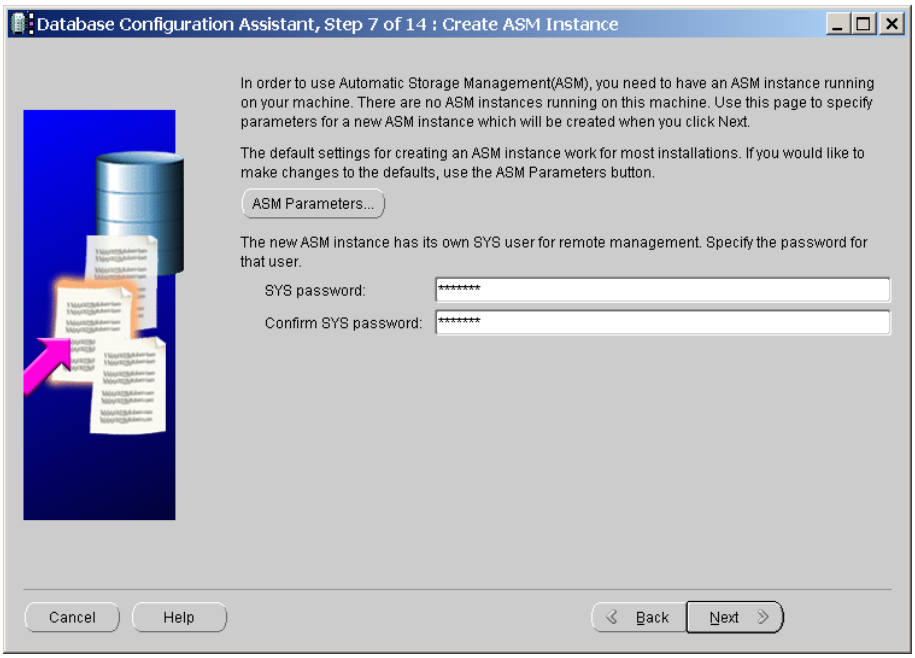

Enter a password for the SYS user of the ASM instance and click "Next".

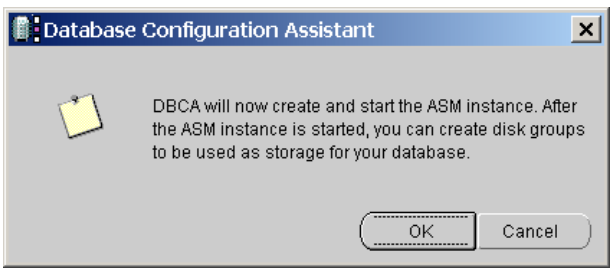

Click OK

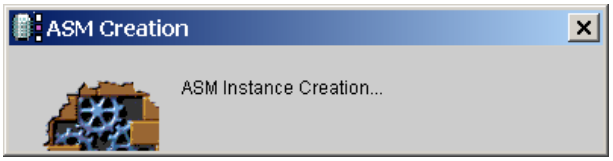

Whilst `ASM Instance Creation` is in progress you will see the above dialog box.

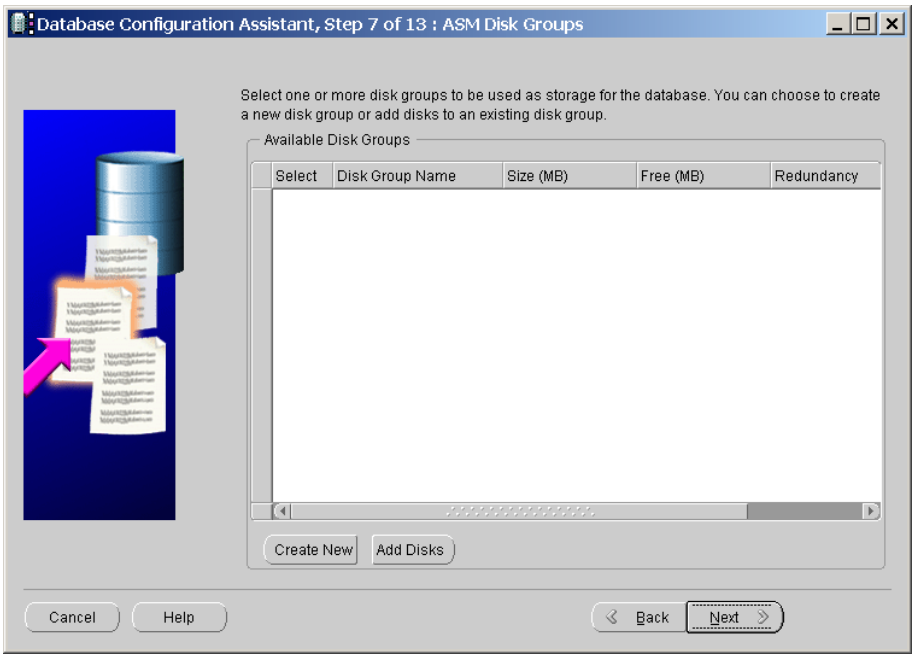

Once theASM instance has been created, a dialog box is brought up which allows us to set up available disk groups. Currently we have no disk groups and we need to create them from new. Click `Create New`.

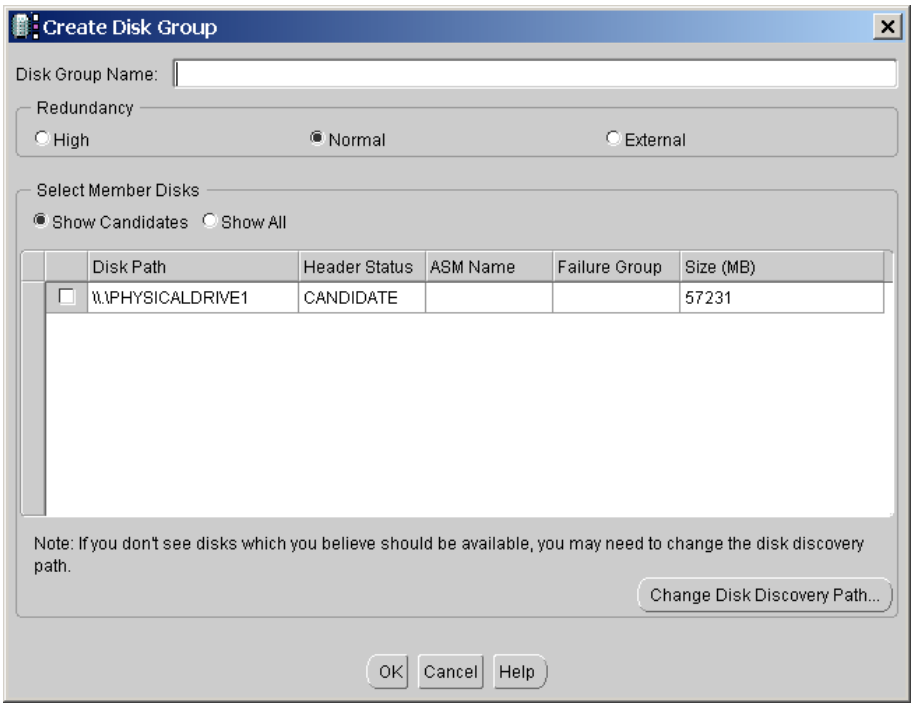

The `Create Disk Group` dialog shows the whole hard drive as the only candidate. To show the logical drives we created earlier we will need to change ASM's disk discovery path. Click `Change Disk Discovery Path`.

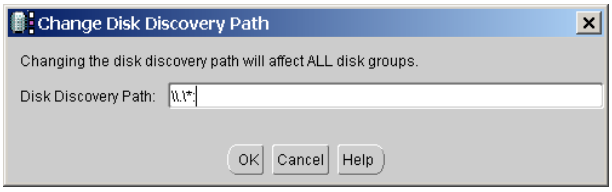

Enter `\\.\*:` and click `OK`.

WARNING: In our example, you must not use D: or C: as these are real disks containing Windows filesystems. To check this you should look in Windows Disk Manager at the Filesystem column (see start of section 2.2).
If you choose real disks you will probably lose the data on them

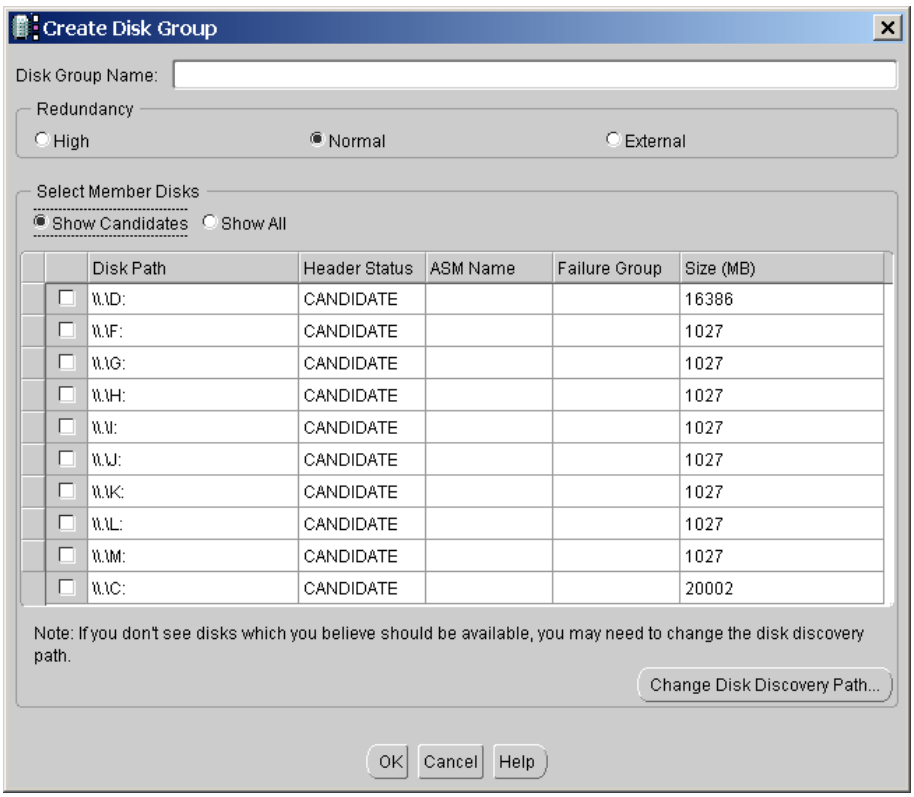

You will now see all the logical drives. NOTE: you may have to select "`Show All`" and then "`Show Candidates`" to refresh the display

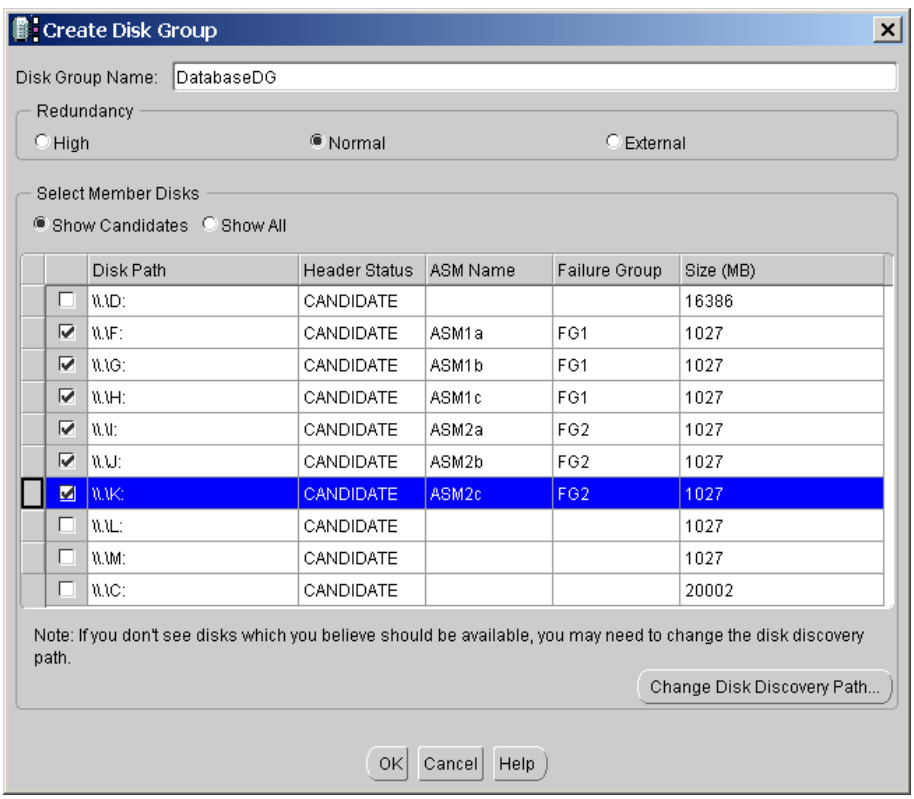

Enter DatabaseDG for the Disk Group Name and ensure that "Normal" redundancy is selected

Select the check boxes next to logical drives F: through K: and enter the ASM names and Failure Group details as in the illustration above

We will create two failure groups FG1 and FG2 and assign three 1Gbyte drives to each. ASM1a, ASM1b and ASM1c will be members of FG1 and ASM2a, ASM2b and ASM2c will be members of FG2. Once complete click OK.

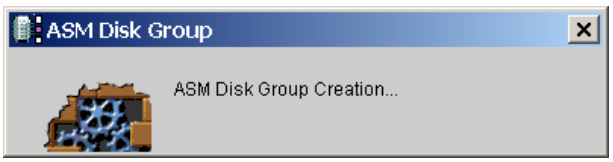

The above dialog box will appear whilst the ASM Disk Group is being created.

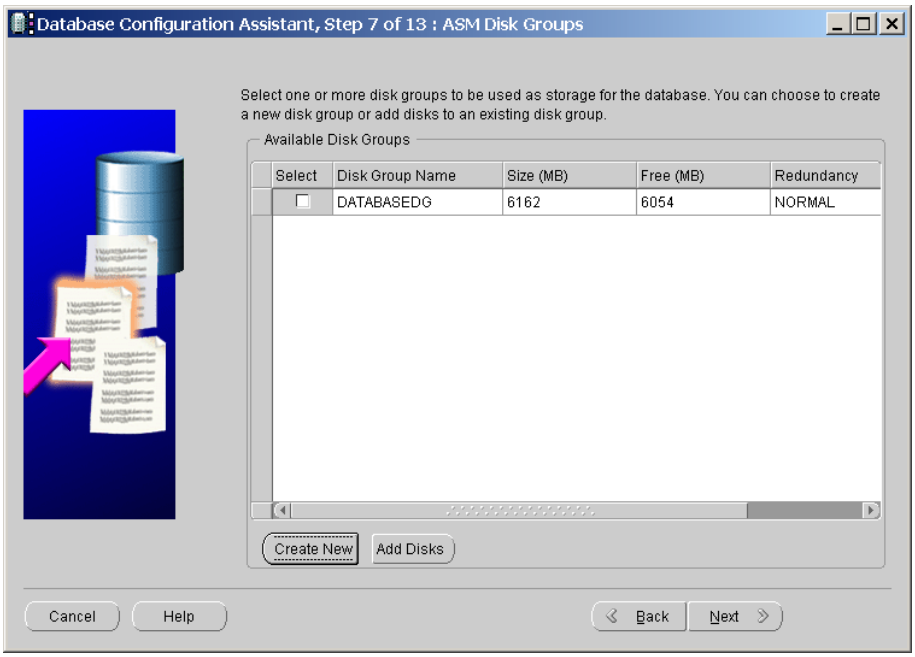

We now wish to create a second disk group for the Flashback Recovery Area.

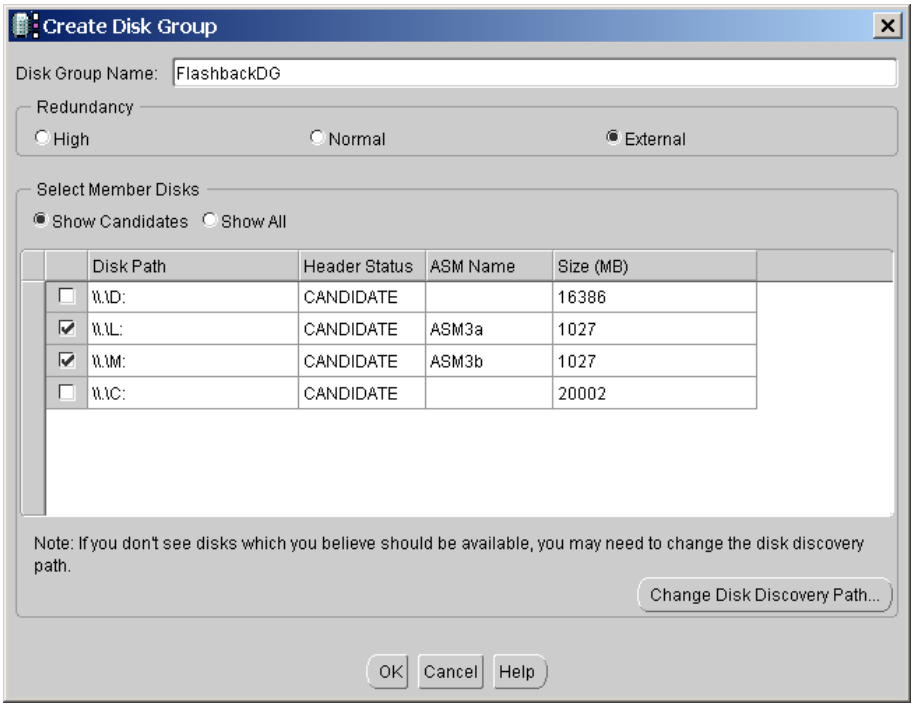

Enter `FlashbackDG` for the Disk Group Name and ensure that "`External`" redundancy is selected, this indicates that we will use external mirroring.

Select the check boxes next to logical drives `L:` and `M:` and enter the ASM names and Failure Group details as in the illustration above. Once complete click `OK`.

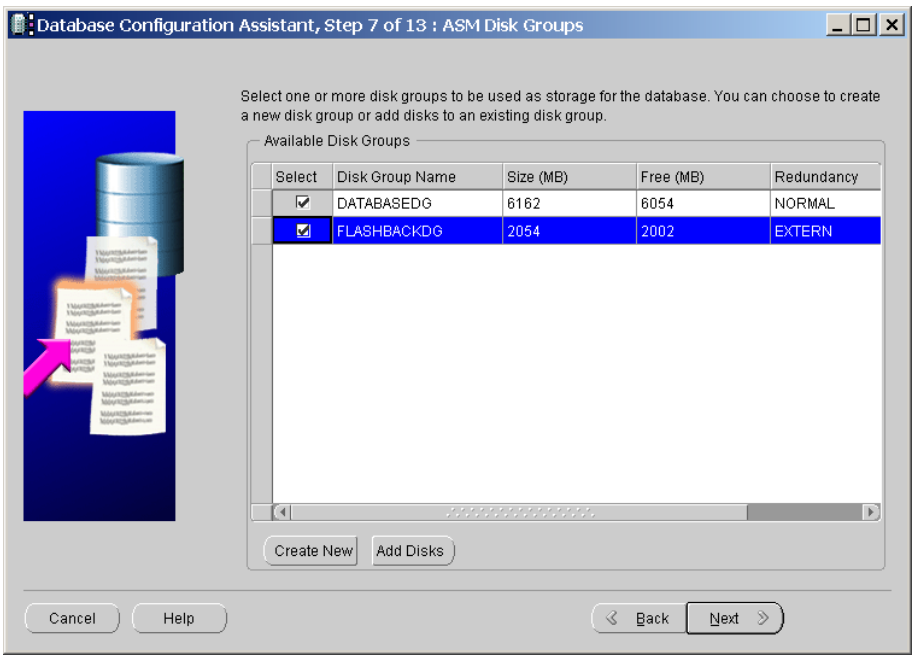

Click "`Next`"

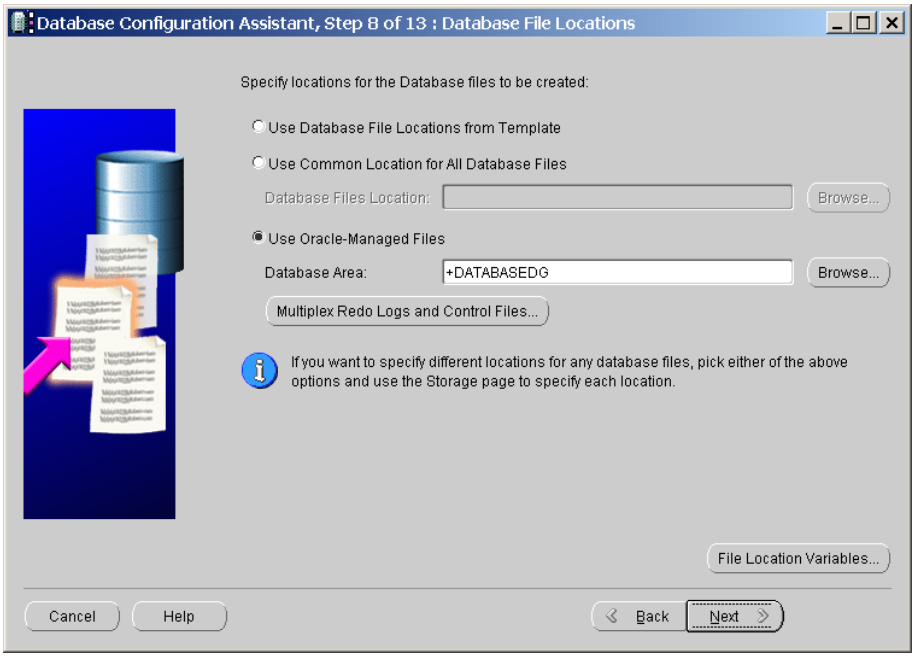

Select "`Use Oracle-Managed Files`" and ensure that the Database Area is set to `+DATABASEDG` then click "`Next`".

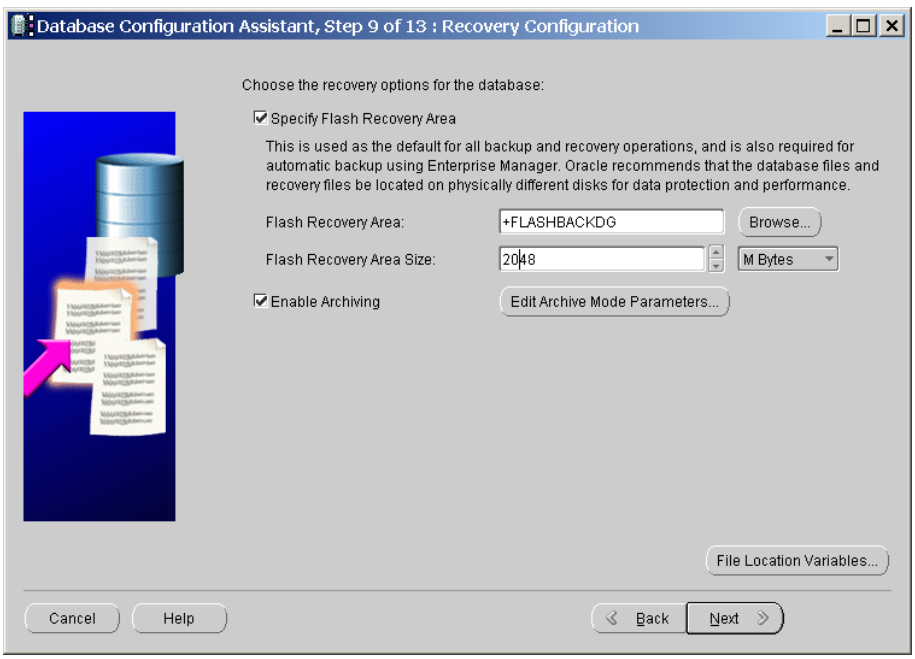

Check "`Specify Flash Recovery Area`", set the name to `+FLASHBACKDG` and set the Size to 2048 M Bytes. The default in the drop down list is Bytes so make sure that M Bytes is chosen. Check "`Enable Archiving`" and click "`Next`".

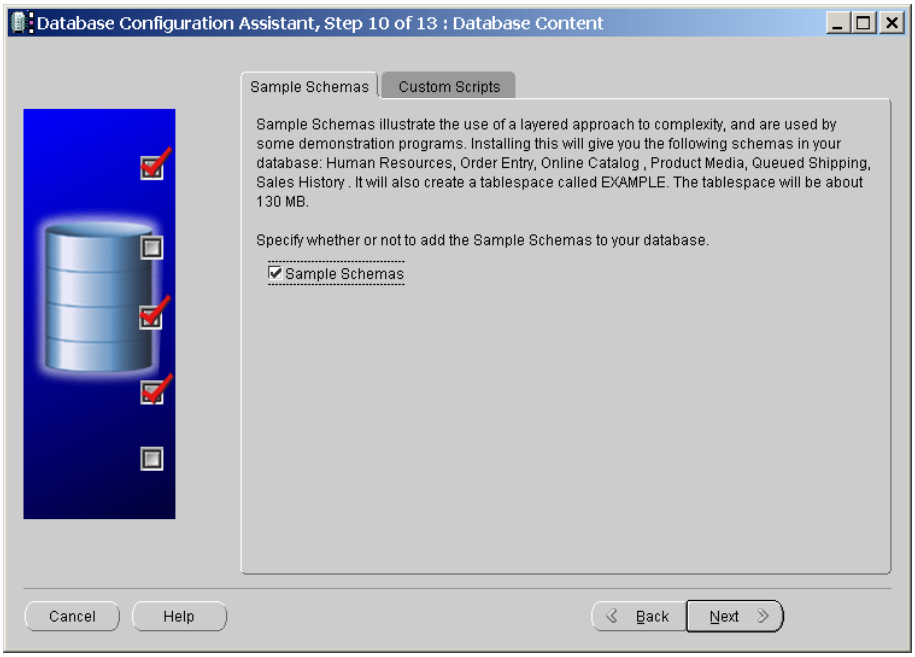

Ensure that "Sample Schemas" is checked and click "Next".

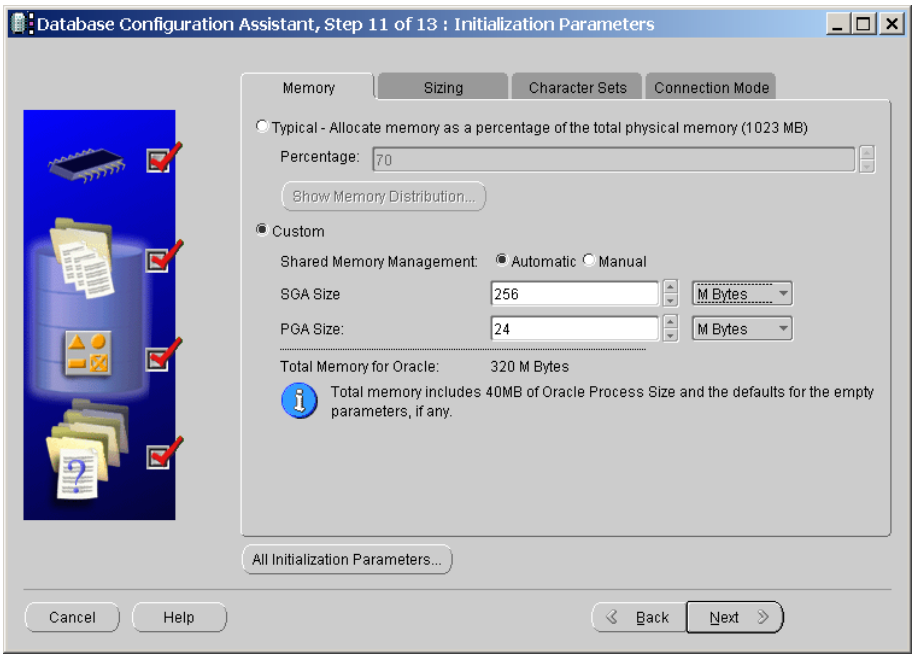

Select "Automatic" for Shared Memory Management and enter an appropriate SGA size in M Bytes. We have chosen 256 M Bytes. Click "Next".

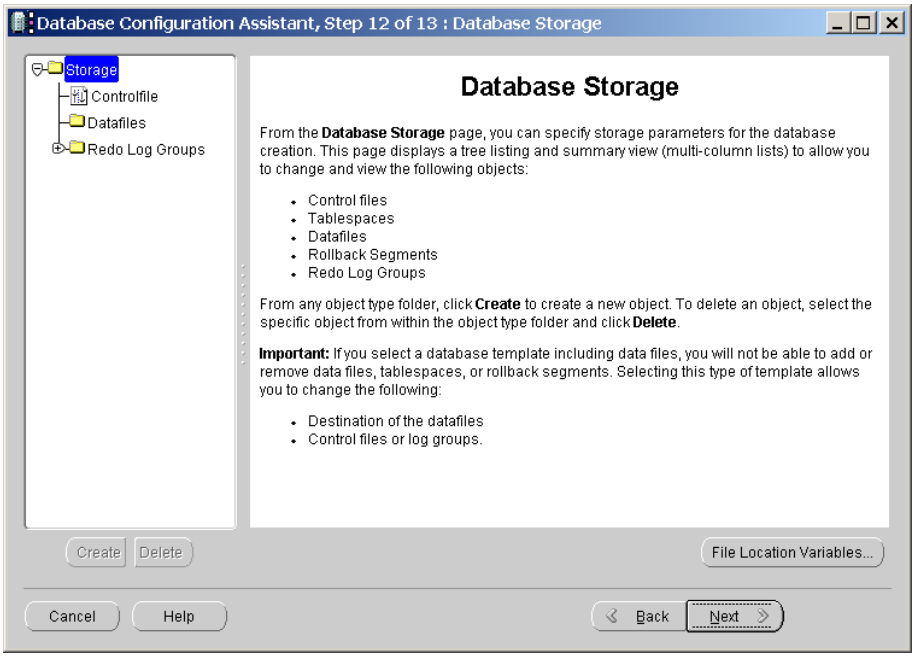

At this point you may be interested in the locations of various files, if you are select the appropriate branches in the tree on the left. Once you are ready to create your database click "`Next`".

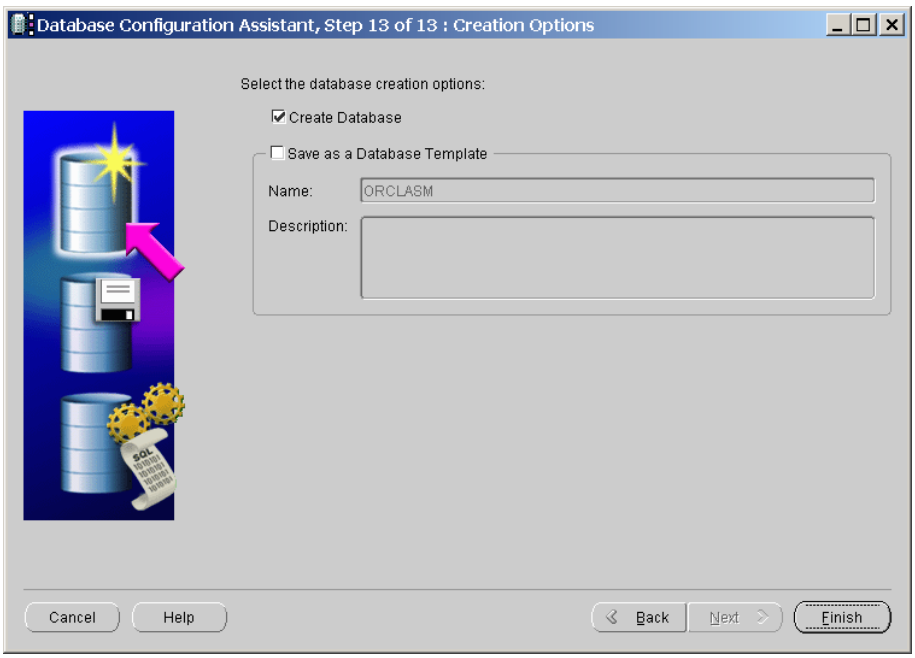

Ensure that "`Create Database`" is checked and click "`Next`".

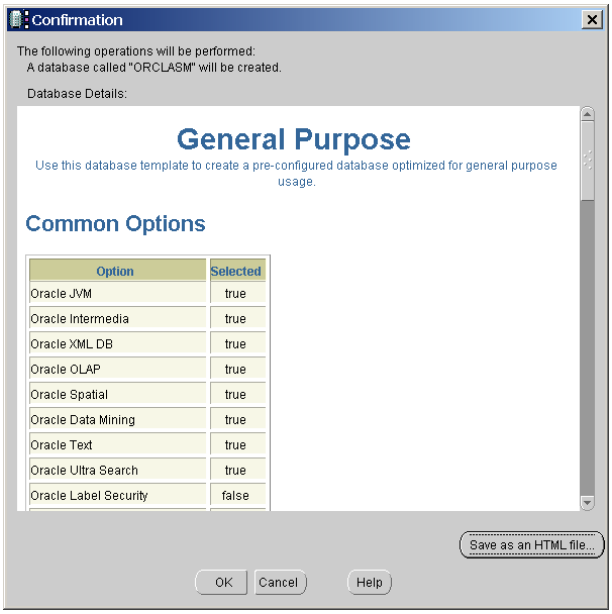

Click OK.

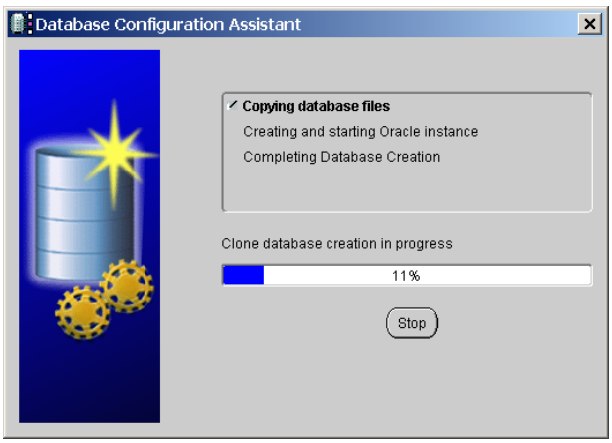

You will see progress from `Copying database files`…

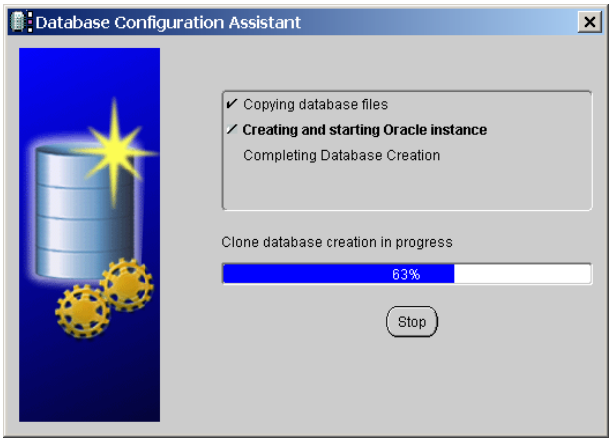

… to `Creating and starting Oracle instance` …

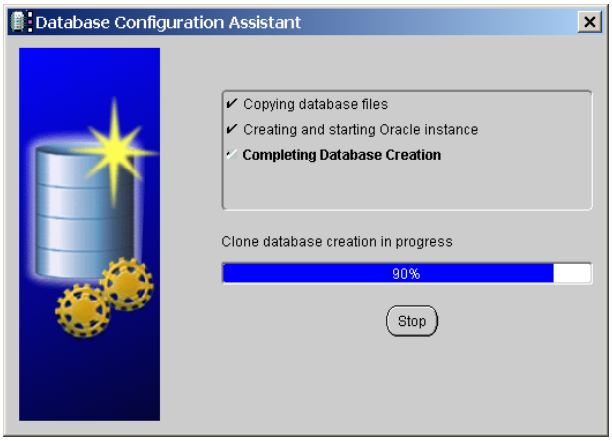

... to `Completing Database Creation`.

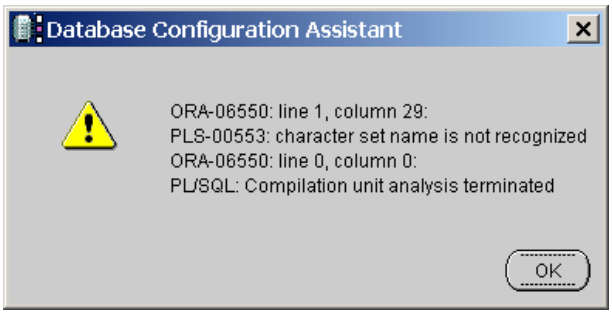

If you see this error, it is a known bug and can be safely ignored. Click `OK`.

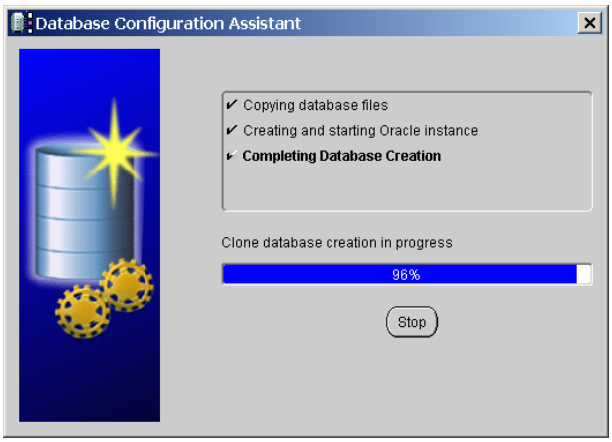

Clone database creation will continue.

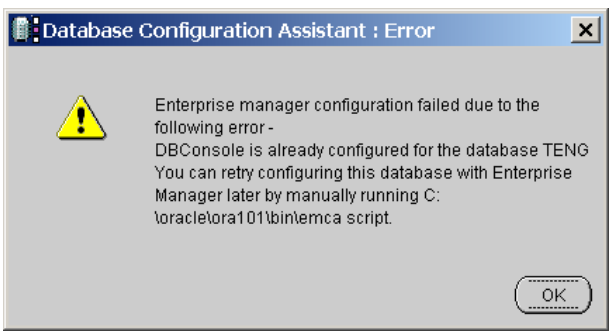

If you see the above message, then you have a previously installed database and you will need to run EMCA –r to allow Enterprise Manager to manage your Oracle ASM database instance.

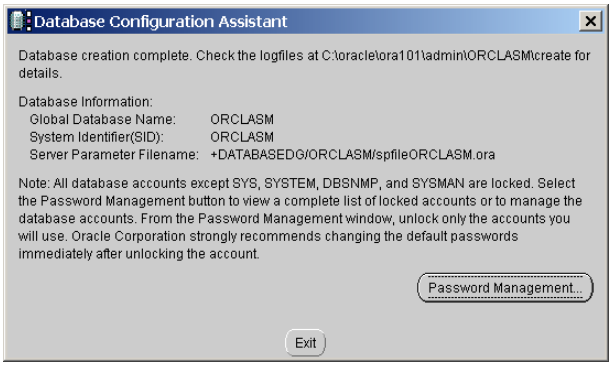

We wish to unlock some accounts for testing purposes. Click `Password Management`.

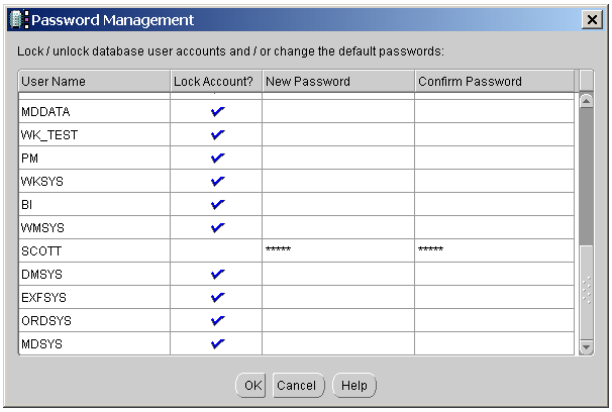

Unlock `HR`, `OE` and `SCOTT` giving them appropriate passwords. Click OK. You now have an Oracle Database `ORCLASM` running with Automatic Storage Management.

4. Hands on One – viewing ORCLASM in Enterprise Manager

If you received an error message regarding support for the DBConsole then try running:

```
emca -r
```

To check which port DBConsole is running on you can use

```
netstat -a
```

If this initially succeeded then the port is probably 5500. Since we had to re-run EMCA our port is 5501. You will be told this port number if you run EMCA and it is also stored in

```
%ORACLE_HOME%\oc4j\j2ee\OC4J_DBConsole\config\http-web-site.xml
```

Launch Internet Explorer from the desktop.

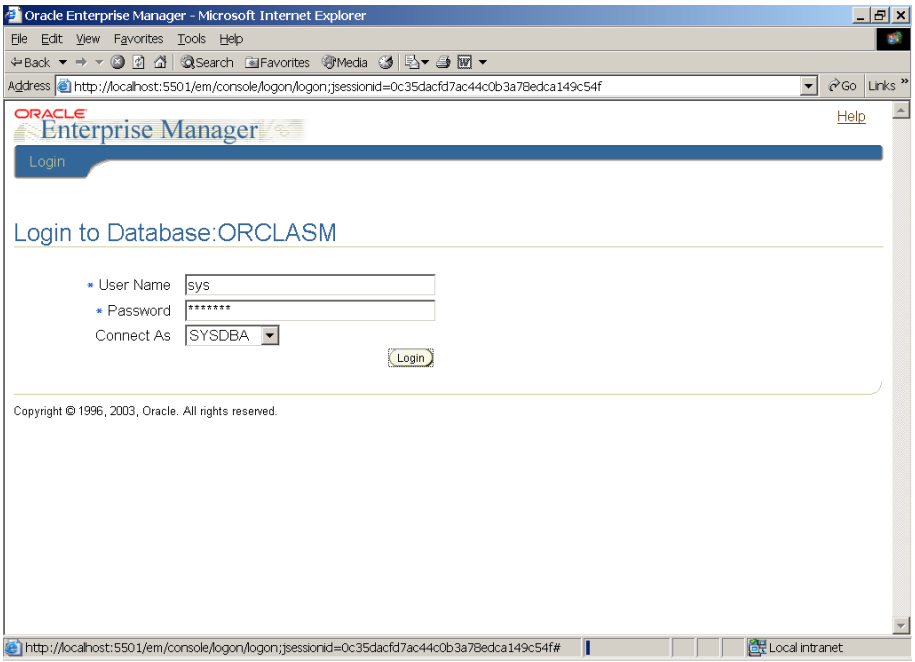

Navigate to http://localhost:5501/em and login as SYS with SYSDBA privileges

Once the DBConsole has swapped into memory you should see the following:

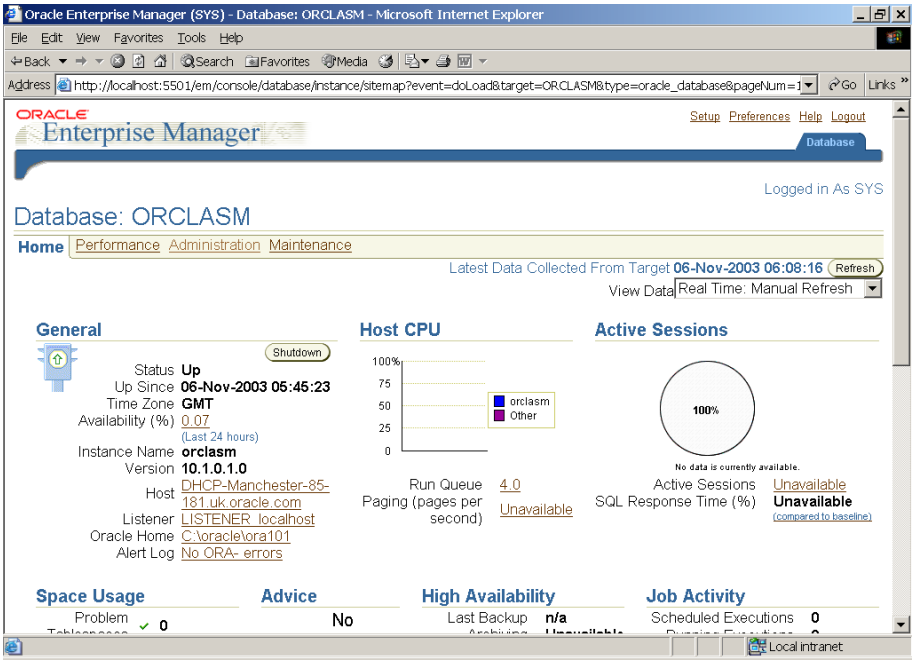

Click `Administration` **è** `Storage` **è** `Datafiles` to see how ASM is used.

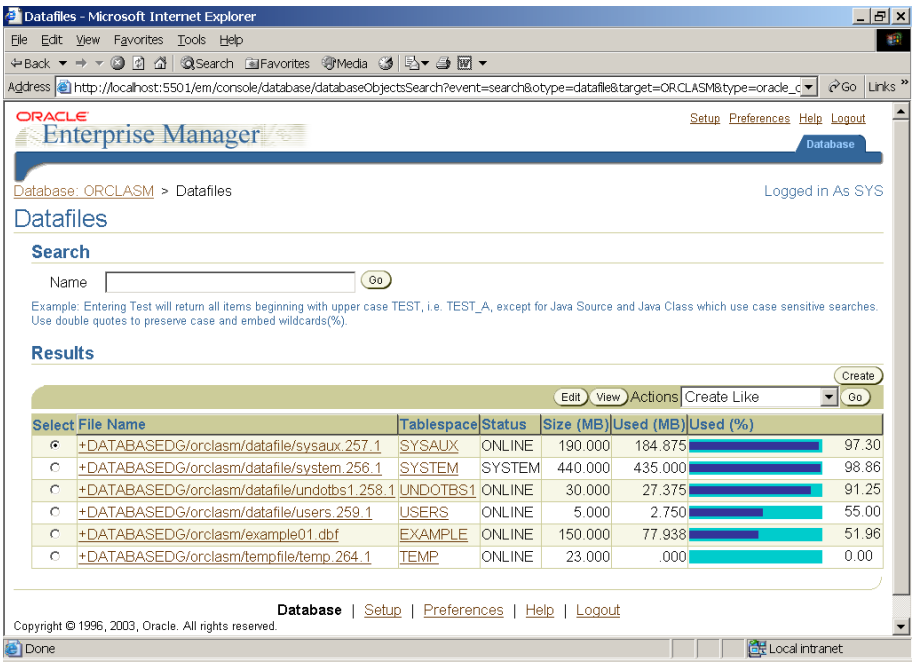

We have already configured Flashback Database so follow “Performing Backups and Recovering Your Database OBE” by selecting Flashback Database and going to step2. For information, you may wish to follow step 1, but it is already set.

Details of this OBE are available at `http://database.us.oracle.com` by selecting the 10g tab and then the OBEs tab.

Congratulations! You have a fully implemented 10g database running on Automatic Storage Management viewable in Enterprise Manager.

5. Hands On Two – Add and Remove Disk Groups. Force a failure.

In deference to all the DBA and Sys Admins out there Hands On Two in contrast to Hands On One is dedicated to those with a command line persuasion.

5.1 Pre-requisites

In order to carry out this Hands On we need two more logical drives. If you are unsure how to do this please refer to section 2.

In our example, we have two further 1Gbyte logical drives labelled `N:` and `O:`

5.2 Add a disk group

To add a disk group we need to manipulate the OSM instance.

From a DOS command prompt:

```
set ORACLE_SID = +OSM
sqlplus / as sysdba
```

Check current settings:

```
SQL> select name, total_mb, failgroup from v$osm_disk where name is not null;
NAME                           TOTAL_MB FAILGROUP
------------------------------ ---------- --------------------
ASM1A                                1027 FG1
ASM1B                                1027 FG1
ASM1C                                1027 FG1
ASM2A                                1027 FG2
ASM2B                                1027 FG2
ASM2C                                1027 FG2
ASM3A                                1027 ASM3A
ASM3B                                1027 ASM3B
8 rows selected.
```

We will add one logical drive to failure group `FG1` with a name of `ASM1D` and one logical drive to failure group `FG2` with a name of `ASM2D`.

```
SQL> alter diskgroup databasedg add failgroup FG1 disk '\\.\N:' NAME ASM1D failgroup FG2 disk
'\\.\O:' NAME ASM2D;
SQL>select name, total_mb, failgroup from v$osm_disk where name is not null order by name;
NAME                           TOTAL_MB FAILGROUP
------------------------------ ---------- --------------------
ASM1A                                1027 FG1
ASM1B                                1027 FG1
ASM1C                                1027 FG1
ASM1D                                1027 FG1
ASM2A                                1027 FG2
ASM2B                                1027 FG2
ASM2C                                1027 FG2
ASM2D                                1027 FG2
ASM3A                                1027 ASM3A
ASM3B                                1027 ASM3B

10 rows selected.
```

You may if you wish check the dump file (`alert_+osm.log`), it should contain entries similar to the following:

```
Thu Nov 06 17:37:15 2003
ORB0 relocating file +DATABASEDG.258.1 (4 entries)
ORB0 relocating file +DATABASEDG.261.1 (2 entries)
ORB0 relocating file +DATABASEDG.262.1 (2 entries)
ORB0 relocating file +DATABASEDG.263.1 (1 entries)
ORB0 relocating file +DATABASEDG.264.1 (3 entries)
ORB0 relocating file +DATABASEDG.265.1 (46 entries)
Thu Nov 06 17:37:36 2003
RBAL stopping process ORB0
RBAL rebalance completed, gn=1
checking for empty OSM disks, gn=1
```

5.3 Drop added disks

We will now drop the disk that we have added

```
SQL>alter diskgroup databasedg drop disk ASM1D,ASM2D;
SQL> select name, total_mb, free_mb,total_mb - free_mb "left", failgroup from v$osm_disk
where name is not null order by name;
NAME                 TOTAL_MB    FREE_MB       left FAILGROUP
-------------------- ---------- ---------- ----------
ASM1A                        1027        701        326 FG1
ASM1B                        1027        708        319 FG1
ASM1C                        1027        714        313 FG1
ASM2A                        1027        696        331 FG2
ASM2B                        1027        714        313 FG2
ASM2C                        1027        713        314 FG2
ASM3A                        1027        928         99 ASM3A
ASM3B                        1027        933         94 ASM3B
```

You may if you wish check the dump file (`alert_+osm.log`), it should contain entries similar to the following:

```
Thu Nov 06 17:18:48 2003
ORB0 relocating file +DATABASEDG.256.1 (23 entries)
ORB0 relocating file +DATABASEDG.257.1 (45 entries)
Thu Nov 06 17:19:14 2003
ORB0 relocating file +DATABASEDG.258.1 (6 entries)
ORB0 relocating file +DATABASEDG.261.1 (1 entries)
ORB0 relocating file +DATABASEDG.262.1 (1 entries)
ORB0 relocating file +DATABASEDG.263.1 (2 entries)
ORB0 relocating file +DATABASEDG.264.1 (3 entries)
ORB0 relocating file +DATABASEDG.265.1 (45 entries)
Thu Nov 06 17:19:37 2003
RBAL stopping process ORB0
RBAL rebalance completed, gn=1
checking for empty OSM disks, gn=1
OSM disk ASM2D successfully emptied
OSM disk ASM1D successfully emptied
```

5.4 Simulate a failure group failure

We will now "kill" failure group FG2 by forcing a drop:

```
SQL> alter diskgroup databasedg drop disks in failgroup FG2 FORCE;
```

You can monitor how this works by running the query from above to see how the remastering etc. is going. Also the logs contain useful information as follows:

```
Thu Nov 06 17:51:40 2003
RBAL CIC release invoked, gn=1, full=0
Thu Nov 06 17:51:40 2003
Determining relations for group DATABASEDG
Determining partners for group DATABASEDG
ORA relocating file +DATABASEDG.1.1 (2 entries)
ORA relocating file +DATABASEDG.2.1 (1 entries)
ORA relocating file +DATABASEDG.3.1 (42 entries)
Thu Nov 06 17:51:41 2003
RBAL re-discovering group 1/0x200AF83 (DATABASEDG), full=0
Thu Nov 06 17:51:41 2003
ORA relocating file +DATABASEDG.4.1 (2 entries)
ORA relocating file +DATABASEDG.5.1 (1 entries)
ORA relocating file +DATABASEDG.6.1 (1 entries)
ORA relocating file +DATABASEDG.256.1 (120 entries)
ORA relocating file +DATABASEDG.256.1 (120 entries)
ORA relocating file +DATABASEDG.256.1 (120 entries)
ORA relocating file +DATABASEDG.256.1 (82 entries)
ORA relocating file +DATABASEDG.257.1 (120 entries)
ORA relocating file +DATABASEDG.257.1 (72 entries)
ORA relocating file +DATABASEDG.258.1 (31 entries)
ORA relocating file +DATABASEDG.259.1 (6 entries)
ORA relocating file +DATABASEDG.260.1 (8 entries)
ORA relocating file +DATABASEDG.261.1 (16 entries)
ORA relocating file +DATABASEDG.262.1 (16 entries)
ORA relocating file +DATABASEDG.263.1 (16 entries)
ORA relocating file +DATABASEDG.264.1 (24 entries)
ORA relocating file +DATABASEDG.265.1 (120 entries)
ORA relocating file +DATABASEDG.265.1 (32 entries)
ORA relocating file +DATABASEDG.266.1 (1 entries)
Thu Nov 06 17:51:42 2003
group DATABASEDG: relocated PST to:
disk 0003 (PST copy 0)
Thu Nov 06 17:51:42 2003
RBAL queued rebalance (power 1) for group 1/0x200AF83 (DATABASEDG)
RBAL CIC discovery already done, gn=1
```

```
RBAL resuming rebalance of group 1/0x200AF83 (DATABASEDG)
Starting background process ORB0
ORB0 started with pid=11, OS id=1908
```

Finally, you will end with only one Failure Group and the database will continue to run:

```
Thu Nov 06 17:52:35 2003
ORB0 relocating file +DATABASEDG.258.1 (16 entries)
ORB0 relocating file +DATABASEDG.259.1 (3 entries)
ORB0 relocating file +DATABASEDG.260.1 (4 entries)
ORB0 relocating file +DATABASEDG.261.1 (8 entries)
ORB0 relocating file +DATABASEDG.262.1 (8 entries)
ORB0 relocating file +DATABASEDG.263.1 (8 entries)
ORB0 relocating file +DATABASEDG.264.1 (12 entries)
ORB0 relocating file +DATABASEDG.265.1 (77 entries)
ORB0 relocating file +DATABASEDG.266.1 (1 entries)
Thu Nov 06 17:52:52 2003
RBAL stopping process ORB0
RBAL rebalance completed, gn=1
Thu Nov 06 17:52:53 2003
group DATABASEDG: relocated PST to:
disk 0003 (PST copy 0)
Thu Nov 06 17:52:53 2003
checking for empty OSM disks, gn=1
OSM disk ASM2A successfully emptied
OSM disk ASM2B successfully emptied
OSM disk ASM2C successfully emptied
```

Congratulations! You have successfully implemented ASM, have added and removed disks and have simulated a failure. All with zero downtime.